A COMPUTATIONAL ANALYSIS OF THE ROBBINS-MONRO STOCHASTIC APPROXIMATION METHOD

THOMAS J. STILLINGS
ROBERT S. LOGAN

A COMPUTATIONAL ANALYSIS OF THE ROBBINS-MONRO STOCHASTIC APPROXIMATION METHOD

by

Thomas J. Stillings

Lieutenant, United States Navy

and

Robert S. Logan

Lieutenant, United States Navy

Submitted in partial fulfillment of
the requirements for the degree of

MASTER OF SCIENCE
IN
OPERATIONS RESEARCH

United States Naval Postgraduate School
Monterey, California

1 9 6 5

A COMPUTATIONAL ANALYSIS OF THE ROBBINS-MONRO STOCHASTIC APPROXIMATION METHOD

* * * * *

Thomas J. Stillings

and

Robert S. Logan

A COMPUTATIONAL ANALYSIS OF THE ROBBINS-MONRO STOCHASTIC APPROXIMATION METHOD

by

Thomas J. Stillings

and

Robert S. Logan

This work is accepted as fulfilling

the thesis requirements for the degree of

MASTER OF SCIENCE

IN

OPERATIONS RESEARCH

from the

United States Naval Postgraduate School

ABSTRACT

The problem of determining extreme percentiles of a distribution function by methods based on the Robbins-Monro procedure is examined by computer simulation. Several techniques for improving the convergence properties are proposed and studied. A stopping rule based on binomial confidence limits is examined.

ACKNOWLEDGEMENTS

We wish to thank Professor Robert R. Read of the Statistics Department, U. S. Naval Postgraduate School, for his encouragement and most helpful suggestions. Lt. Logan was largely responsible for Section C, sub-sections 2, 3 and 4 and Lt. Stillings for sub-sections 5 and 6. Sub-section 7 and all other facets were joint efforts.

TABLE OF CONTENTS

A. The Problem

Consider a problem in sensitivity testing where items are selected at random from a population and subjected to a known stimulus level x. The actual, critical stimulus level of each item, X, is a random variable whose distribution $M_X(x)$ is the distribution of the critical stimulus levels in the population. X cannot be directly observed but information about it is gained by observing the response or non-response of each item when the known level x is applied. The magnitude of x may vary from item to item in the sample and may depend, at each trial, upon previous observed outcomes. Assume, also, that each item within the sample is subjected to only one magnitude x and is tested only one time.

Let $Y(x)$ be a random function such that

$$(1)\quad Y(x) = \begin{cases} 1 \text{ if } X \leq x & \text{(response occurs)} \\ 0 \text{ if } X > x & \text{(non-response occurs)} \end{cases}$$

Therefore, to each value x there corresponds a random variable $Y(x)$ with a distribution function $P\left[Y(x) \leq y\right]$ such that if $P\left[Y(x) \leq y\right] = G_Y(y/x)$ then $M_X(x) = \int_{-\infty}^{\infty} y\,dG_Y(y/x) = P(X \leq x) = E\left[Y(x)\right]$

Let α, $0 < \alpha < 1$, be a given constant such that the equation $M_X(x) = \alpha$ has a unique but unknown root, θ, i.e., $M_X(\theta) = \alpha$. We are interested in estimating this root where α is the proportion of the population which will respond to stimulus level θ.

Parametric and non-parametric methods exist for the class of problems outlined above. However, known methods have not been found practical in the case of sensitivity testing for an extreme percentile of a distribution. The parametric methods such as (5) currently used by NOL fail because of the drastic limitations the assumed distribution imposes on the tail of

the function $M_X(x)$. The non-parametric methods seem to require extremely large sample sizes and lack a workable stopping rule, i.e., a technique to tell us when to stop sampling while giving us some "measure of closeness" to the true θ point.

Among the most popular of the non-parametric methods is the Robbins-Monro stochastic approximation method. The purpose of this paper will be to computationally, rather than analytically, investigate this method in an effort to: first, find ways of improving its convergence; second, outline the method for a stopping rule.

B. Background

The non-parametric stochastic approximation method suggested by Robbins and Monro (1) involved selecting a sequence of positive constants a_n , $n \geq 1$, such that

(2) $\sum_{n=1}^{\infty} a_n = \infty$ and $\sum_{n=1}^{\infty} a_n^2 < \infty$ and constructing a non-stationary Markov chain $\{x_n\}$ to serve as the approximation of θ where

(3) $x_n = x_{n-1} - a_n \left[Y_n (x_{n-1}) - \alpha \right]$, $n \geq 1$, and the zeroth approximation, x_o, is an arbitrary gures for θ . They proved that this sequence of point estimates of θ will converge to the true θ point in probability. It should be noted that M(x) as used in their paper, referred to a much more general class of functions of which the continuous probability distribution has been proven to be a subclass. Blum (2), using two additional assumptions on M(x), proved that x_n will converge to θ with probability 1.

The main questions associated with this method are the speed of convergence and the choice of the coefficients, a_n, to maximize the speed. Chung (3) attacked this problem by studying the asymptotic behavior of the moments of X_n and was able to prove asymptotic normality under certain conditions. In particular, theorem 9 (3) with $a_n = c/n$, then $n^{1/2} (X_n - \theta)$ tends to the normal distribution $N \left[0, \sigma^2 c^2/(2\alpha_1 c-1)\right.$ in probability, where $\alpha_1 = M'(\theta) > 0$ and σ^2 is the variance of $Y(\theta)$. Thus the asymptotic variance of X_n tends to zero with the speed 1/n. But among the assumptions in this theorem Chung notes that

$\lim_{|x| \to \infty} \left[M(x)/x \right] > 0$

is very restricted from a statistical point of view since M(x) is not, in this case, a bounded function and therefore not a distribution.

Chung studies another case, theorem 6 (3), in which M(x) is bounded, where the coefficients are $a_n = 1/n^{(1-E)}$, where E must exceed a positive number $1/2(1+K_4)$ and where K_4 depends on the problem. He shows that $n^{(1-E)/2}(X_n - \theta)$ will have a normal limit so that the variance of X_n tends to zero with the speed $1/n^{(1-E)}$. But here also it is pointed out that this speed is not statistically practical.

In an effort to resolve the above problem Hodges and Lehmann (4), using theorem 1 (2), prove that theorem 9 (3) remains valid when the

condition $\lim_{|x| \to \infty} \left[M(x) / x \right] > 0$

is removed thus allowing the result to be used when M(x) is bounded, e.g., continuous distribution functions are therefore allowable. In particular, they recommend the use of $a_n \sim c/n$ and suggest that c be chosen so as to minimize the asymptotic normal variance

$\sigma^2 c^2/(2\alpha_1 c-1)$. It can be shown that $c = 1/\alpha_1$, is a minimizing value of this variance and will reduce it to the form $(\sigma / \alpha_1)^2$. However, it is pointed out that in practice α_1, is not itself known and must be guessed. Since it is assumed that $M'(\theta) = \alpha_1 > 0$, it is suggested that a "safe" small a priori estimate for α_1 be used to make $c > 1/2\alpha_1$.

C. Methods of Attack

1. Introduction

Any practical application of the Robbins-Monro process, i.e., to obtain a reasonably accurate approximation of θ with a limited sample size, forces us to make two assumptions which are not impossible to fulfill in areas where use of the process might be considered. The first is that the experimenter be capable of drawing large sample sizes, $n > 200$. The second is that enough previous data including, possibly, past parametric results should be available to enable the experimenter to make an initial guess, x_0 , which is reasonably close to the true point. In most of the methods of attack considered in this paper it will be tacitly assumed the above conditions are obtainable.

The authors will also follow the recommendation in (4) and use the coefficients $a_n = c/n$, where c is a constant with the addition of a case where c is a function of n and, of course, equation (2) must be satisfied. Therefore, we are now employing the recursion relation,

$$(4) \quad x_n = x_{n-1} - c/n \left[Y(x_n) - \alpha \right], \quad n \geq 1$$

where $Y(x_n)$ is defined as in equation (1). Also we are restricting our attention to large values of α since the character of the problem and the results may be complemented to apply to the case of small values of α .

The mathematical analysis of this problem is extremely difficult and for this reason our mode of attack is via computer simulation in the hope that some practical quantitative information can be obtained to serve as a guide for future analyses.

The first difficulty to emerge is that, for large α , movement of the approximation, x_n , to the right far outbalances movements to the left, i.e., from equation (4) movement is either $c/n(\alpha)$ or $-c/n(1-\alpha)$ to the right or left, respectively. As a result even moderate deviations from the expected success-failure pattern may drive the approximation, x_n , so far from θ that recovery becomes increasingly difficult as n grows large. The likelihood of this occurrence can be diminished by using either a "moderate" value of c or a deliberately high value for the guess, x_o .

The effect of the choice of c is discussed in sub-section 2 while the use of a deliberately large x_o is treated in sub-section 3. The use of a priori barriers on the x_n is considered in sub-section 4. A comparison of sample variances with the theoretical asymptotic variance is made in sub-section 5 and a block averaging technique is taken up in sub-section 6. An endeavor, which we call a pseudo confidence interval, to obtain a stopping rule utilizing the success-failure pattern of the process is outlined in sub-section 7.

2. The Value of c

As noted in the introduction, too large a value of c leads to great oscillation in the $\{x_n\}$ while too small a value results in insufficient movement. Due to the impossibility of finding an optimal c, which would depend upon the actual distribution and the accuracy of the guess both of which are unknown, we have left its determination up to the overall results of the computational analysis. Only a cursory preliminary investigation was made in an effort to find a "range of c" for the two diverse distributions, the standard normal and the uniform on the interval (0,1). That is, assuming knowledge that those were the

actual distributions of the critical stimulus levels, equation (4) was utilized for various values of x_o and c while, at each stage, comparing the value of the current x_n with the true θ point and stopping if the difference between the two was less than or equal to some small value. Overall results are listed in Appendix 1 while, for the reader's immediate information, a summary of the "best" (i.e. $|X_n - \theta|$ 0.001)c for various x_o is given here.

U(0,1)	c	x_o	n	α
	0.6-1.0	0.8	25-150	.95
	0.7-1.0	0.8	25-100	.96
	0.9-1.0	0.8	50-100	.97
	0.8-1.0	0.8	25-100	.98
	0.5-1.0	0.9	25-50	.95
	0.9-1.0	0.7	25	.95
	1.0	0.6	20	.95

N(0,1)	c	x_o	n
α =0.95	1.6-3.0	1.0	50-400
	1.7-6.5	1.2	20-150
	1.8-4.5	1.5	10-350
	1.8-4.0	2.0	10-250
	3.5-6.5	2.5	25-150

The values of n above are approximately the range for the various values of c at which ($x_n - \theta$) $\leq$ 0.001.

The accuracy of the guess x_o seems to be the most sensitive parameter, i.e., for close guesses moderate values of c seem to be best while for more distant guesses larger values of c work best. Further, with a distribution such as the normal overall results seem to be heavily dependent upon which side of θ the guess is located.

The asymptotic analysis of Chung (3) seems to indicate the use of a larger value for c than is recommended by our results. However, it must be emphasized that Chung was studying theory which would hold uniformly and, hence, for very distant guesses; and that a priori information and limited sample size was not a consideration.

3. The High Guess

With the observation about the amount of movement of the sequence given in the introduction it seems that a high guess would "most probably" be better than a low one. Our results seem to indicate that it is better.

Therefore, a deliberately high guess could be used to lessen the likelihood that the "best" guess is below the θ point. However, this should be done cautiously. For example, using the N(0,1) as the test distribution and α = 0.95 with c = 2.0-3.0, any high guess up to about 2.0 will generally result in a sequence that arrives in the neighborhood of the true θ point in about 200 iterations. However, with guesses above this point moderate values of c are overwhelmed by n so quickly that x_n seldom gets close to θ .

As α increases,the value of a high guess is even more important. Data seems to indicate that for a low guess, with α = 0.99, the result after n iterations will be farther away from θ than would a correspondingly low guess with a smaller α .

Any decision to increase a "best" guess by some increment before starting the process should be made using past data and results as a basis.

4. A Priori Static Barriers

Since the risks involved with high or low guesses have been pointed out in previous sub-sections, we feel it only necessary to recall the first two assumptions made in the introduction to this section to proceed with a discussion of this method.

Given these assumptions, it seems reasonable to assume an experimenter who, confident in the accuracy of his initial guess, may also be confident enough to place a priori limits around this guess beyond which the x_n cannot travel. The reason for wanting limits on x_n is obviously to avoid the risks pointed out earlier. Further, the existence of limits tends to make the choice of a value for c much less critical.

Since this was not the only method studied,only one form of this method was programmed. Barriers were set symmetrically about our best guess and whenever x_n was either greater than the upper bound or less than the lower bound it was set equal to x_{n-1}. The iterations were allowed to continue but n was regarded as having been increased and a new $Y(x_n)$ was used, i.e., a new success-failure comparison was made using a new random variable. In other words, given an a priori interval, the experimenter simply disregards certain data points if they violate this interval. Thus, the barriers possess a reflecting quality and the number of these reflections are counted. The results are listed in Appendix 2 and the method seems to be promising given θ is within the barriers.

It is obvious that the x_n would tend to reflect on the lower barrier more than on the upper one even if θ is contained within the interval. However, if θ is not in the interval, one expects the sequence to reflect upon one of the barriers an "abnormal" number of times and/or the x_n to be driven close to this barrier. It should be possible to develop a rule for statistically testing this contingency. However, in this case, there is a difficulty in that the cause of inordinate numbers of reflections may be the result of too large a value of c.

The values of c used cover a fairly wide range from too low a value, when no reflections occurred, to high values, when one wastes samples reflecting back and forth too often. These reflections will obviously occur in the first few iterations. But this is exactly the reason the experimenter set the bounds.

It is still important that c should be within some "good" range of values since the investigator wants the process to settle down quickly yet possess enough reflections so that, after all computation is completed, their number and the final value of x_n can give him some indication of the accuracy of his guess and/or the validity of his a priori interval. We recommend a value of c be calculated only after the length of the interval has been set and should be of such a magnitude as to allow one or two reflections on the upper barrier in the first few iterations given a low probability failure occurs.

It should be noted that if the investigator has indications that the interval was not correct he can simply widen the barriers and, if desired, shift his guess and recompute the x_n. If the barriers were simply widened with no concomitant shifting of the guess the x_n computed before

the reflection are still valid and the process can be started over at this point.

There are many variations of this method. Another variation which might be empirically feasible, is some sort of "squeeze down" method in which just enough iterations are made to allow the process to settle down, i.e., where n is so large that reflections cannot possibly occur. Then the experimenter can stop and, after noting the number of reflections on either barrier and his current value of x_n, decide to continue the process or shift the barriers closer about his current x_n, using it as his new "best guess", recomputing a new value for c and starting the process again, i.e., with n = 1 but with different random samples. This is admittedly dangerous since it may be argued that even if Θ were within the interval at the beginning of the process that, when the barriers are shifted around x_n and are made closer by some incremental value, one is assuming that: (a) the x_n is actually closer to Θ than was x_o and (b) the investigator is in possession of enough knowledge that he will not reduce the length of the interval to such an extent that Θ will be outside the interval.

At any rate, the authors believe the general method is advantageous by lessening the harmful effect of large oscillations and/or low guesses and should be used if possible.

5. Comparison of Sample Variance with Asymptotic Variance

In an attempt to compare the behavior of the process for moderately large sample sizes with the theoretical asymptotic variance., the process was carried out to x_{300} twenty five times. For each x_n, 25-sample means and variances were computed. According to the asymptotic theory, $\sqrt{n}(x_n - \theta) \rightarrow N(0, \frac{\sigma^2 c}{2\alpha_1 - 1})$ and one might expect the sequence of 25-sample variances of x_n eventually to behave like $\frac{1}{n} \frac{\sigma^2 c}{2\alpha_1 - 1}$ (although it has not been proven that this is necessarily the case). The attempted comparisons were very disappointing.

A data analysis showed the sample variance to be decreasing as the logarithm of the n^{th} trial. A least squares fit of s^2 vs $\ln(n)$ was computed, where

$$s^2 = 1/m \sum_{j=1}^{m} (x_{ij} - \bar{x}_i)^2$$

$$\bar{x}_i = 1/m \sum_{j=1}^{m} x_{ij}$$

i = the i^{th} trial

m = number of replications of the experiment.

The fit was very good to the 2nd decimal. The mean and variance are also highly insensitive to either x_o or c. The results were,

$S^2 = - B \ln(i) + A$ where

$A \approx 3.9, \quad B \approx 0.42$

It should be noted that the variance must be positive, $\ln(i)$ becomes infinite with i and the intercept is in the neighborhood of i = 1200.

These results and a comparison with the asymptotic variance appear in Appendix 3. Clearly, the process is not at all close to the asymptotic values and yet it has become relatively stable.

6. Block Averaging

The importance of a good initial guess and the inactivity of the sequence when n becomes large has been emphasized throughout this paper. The latter fact, in particular, suggests the following technique might be of value.

Given x_o choose some "large" value of c and some "suitable" block size, m. Experiment to find x_m as usual and call this value $x_{m,1}$. Then use this value to initiate the sequence again; i.e., letting $x_{m,1}$ be the starting point, proceed to generate a second block of m values using a decreased value of c. Let $x_{m,2}$ denote the "new" x_m and average $x_{m,1}$ and $x_{m,2}$ and obtain the initial "guess" for a third block of size m, namely $x_{m,2}$. In general, c is reduced in each block and the initial value for the j + 1th block is

$$x_{o,j+1} = x_{m,j}$$

Iterate in this manner for k blocks so that km = n is the total available sample size.

Two cases were simulated and the results appear in Appendix 4; one for a low and one for a high guess, c is halved before beginning each block, m = 20, k = 10. Another case using c = ln 20 in the initial block while the c in later blocks utilized the logarithm of the value of c immediately preceding. Only positive values of the constant were allowed.

This method seems to have some merit although three difficulties immediately observable are: (a) the value of c to be used initially,

(b) the way this value should be decreased, and (c) the determination of the size of each block.

7. A Pseudo Confidence Interval Stopping Rule

Assuming we start in the vicinity of θ and the process does not "stray" too far, we may view each $Y(x_i)$ as a Bernoulli trial at level, approximately, α. In this case we may place "confidence" limits ($\underline{P}$, $\overline{P}$) on α, i.e. $\text{CONF}(\underline{P} \leq \alpha \leq \overline{P}) = 1 - \alpha_2$ where $\overline{P}$ is the solution to the equation

$$(5) \quad g(\overline{P}) = \sum_{k=x+1}^{n} \binom{n}{k} P^k (1-P)^{n-k} - 1 + \alpha_2/2 = 0$$

and $\underline{P}$ is the solution to the equation

$$(6) \quad g(\underline{P}) = \sum_{k=x}^{n} \binom{n}{k} P^k (1-P)^{n-k} - \alpha_2/2 = 0$$

where n = the number of iterations

x = the number of successes (responses) in n iterations

$(1-\alpha_2)$ = "confidence level" to be achieved before the process is terminated.

For the case $n = x$,

$$\overline{P} = 1, \qquad \underline{P} = (\alpha_2/2)^{1/n}$$

and if $x = 0$,

$$\overline{P} = 1 - (\alpha_2/2)^{1/n}, \qquad \underline{P} = 0.$$

For all other cases the solution to equations (5) and (6) are obtained by use of the recursive formula

$$(7) \quad P_{m+1} = P_m - g(P_m) \,/\, K(n,x)$$

where the initial P_m are the previous solutions for $\overline{P}$ and $\underline{P}$ while the new solutions are obtained at that P_{m+1} such that $g(P_m) \doteq 0$ (i.e., $g(P_m) \leq 0.0001$). By Stirling's approximation for $n!$ it can be shown that $K(n,x) > g'(P_m)$ and, for equation (5),

$$\overline{K}(n,x) = n \left[(n-1)/2\pi\, x(n-1-x)\right]^{1/2} \exp\left[1/12(n-1)\right]$$

for $1 \leq x \leq n-2$, and $\overline{K}(n,x) = n$ for $x = 0, n-1$.

For equation (6),

$$\underline{K}(n,x) = n \left[(n-1)/2\pi\,(x-1)(n-x)\right]^{1/2} \exp\left[1/12(n-1)\right]$$

for $2 \leq x \leq n-1$, and $\underline{K}(n,x) = n$ for $x = 1, n$.

The nesting procedure developed in Appendix 5 may be used to evaluate the g(Pm). The calculations are simple if done on a digital computer where some language such as SCRAP is available.

The sampling continues until the difference $\overline{P} - \underline{P}$ is less than or equal to some preassigned low number. The corresponding x_n is taken as the final estimate of θ.

It seems reasonable to assume that, if we use the preceding "dymamic barrier" technique, there may exist times when a failure causes $\overline{P}$ to go below α or a success causes $\underline{P}$ to go above α. If all trials were taken in a sufficiently small vicinity of θ then either of these contingencies can be interpreted as resulting from a relatively unlikely sample sequence. Thus we reverse the outcomes, i.e., in the former case we change the failure to a success or, in the latter, a success to a failure. In this way the inequality $\underline{P} \leq \alpha \leq \overline{P}$ is always maintained until the stopping point is reached. Of course, x_n, $\overline{P}$ and $\underline{P}$ are re-calculated with the change.

There is a serious danger in such arbitrary changes of data as may be emphasized by considering an extreme case. Suppose x_n is actually a long way above θ , say 10 standard deviations. The method will artificially change the proportion of successes so as to maintain $\underline{P} \leq \alpha \leq \overline{P}$ and the process is prevented from moving in its natural direction. Thus one must start close to θ and use a small value of c. "How close" and "how small a c" are critical questions.

The results of simulation using the regular normal distribution as the sample set appear in 2 groups. The first is a collection of pilot runs in which the simulations were terminated for various programming reasons not related to the stopping rule. However, the information gained was felt to be of value and is, therefore, included in Appendix 6A.

The second set is composed of cases where longer running times were allowed in an effort to analyze the movement of the process and the effect of the stopping rule. In all of these cases $x_o = 2, \alpha = 0.95, \theta = 1.645$ and the analyses of the results are given in Appendix 6B. The system never reached its programmed stopping point, i.e., $\overline{P} - \underline{P} \leq 0.02$. For the several values of c listed it appears that the reversal of outcomes was undesirable in the fashion mentioned previously. That is, our initial value for x_o(0.35 standard deviations above θ) was too far away and, in each case a high number of successes, especially near the beginning of the process, caused $\underline{P}$ to attempt to go above α . In these cases, arbitrarily changing the outcomes and recomputing the x_n drove the process farther away from θ . However, it can easily be seen that, being above θ , if an inordinate number of low probability failures had occurred, an attempt by $\overline{P}$ to go below α would have aided convergence of the x_n when the failure was changed to a success.

It will be noted that c = ln n was utilized in a few cases. This

function for c seems to be of help in certain cases and a proof of the convergence of $\sum_{1}^{\infty} a_n^2 = \sum_{1}^{\infty} (\ln n/n)^2$ is given in Appendix 7.

D. Conclusions

1. The authors believe that the asymptotic methods suggested in the various references while theoretically attractive have little practical value since, as seen from the overall results, the x_n do not change significantly enough when n grows large to be of statistical value. In other words the process will normally settle on one side or the other by at least $100 < n < 200$ while further movement is too insignificant to warrant more samples.

2. The method of a priori barriers on the guess seems to hold significant practical promise. The barriers do keep the x_n "close" to the true θ point and reduce sample size for the same accuracy of x_n given the interval contains the θ point.

3. It is felt that the accuracy of the initial guess, x_0, is the most critical parameter. This is true in all cases, no matter which method is considered, but especially, of course, if barriers are desired.

4. The value of a guess which is "slightly" high cannot be over emphasized, especially when α is large. The higher the value of α the more concerned the experimenter should be to start the process above θ since it is that much more probable that any guess will be driven downward initially when the activity of the process is at its highest level.

5. The attempt to obtain a stopping rule by use of the pseudo confidence interval methods, while failing in a practical sense, nevertheless gives an experimenter one bit of information within $100 < n < 200$ samples. We think it is very important, for reasons previously outlined, to know which side of θ the x_n are located. With the pseudo confidence interval method this bit of information is obtained, within a practical sample size, whenever $\overline{P}$ tries to go below or $\underline{P}$ to go above α . In the

first case the investigator will realize the x_n is probably below θ ; in the second, it is probably above θ . We believe this method, if used with another method which improves convergence, can be used as a rule to "stop when the first reflection occurs" even though no idea of "closeness" is obtained. We also believe that these "violations" of α can be obtained earlier in the process if a larger value of α_2 is used.

In other words, the method of a priori static barriers discussed in subsection 4, Section C, seems to help significantly given sufficient reason for their placement. The pseudo confidence interval method could give this reason for if, after n has reached a sufficiently high value, $\underline{P}$ tried to go above α the experimentation could cease and an upper bound could be placed either at the current x_n point or slightly above it and the process could be re-started, i.e., $x_n = x_o$, drawing further samples. Thus the investigator has some "confidence" for placing this bound on the $\{X_n\}$. In a similar fashion a lower bound could be obtained when $\overline{P}$ tried to go below α .

It is obvious that a method such as this would require quite large sample sizes but they may still be an order of magnitude or so less than that which seems to be required of the unaltered Robbins-Monro technique.

In any case, lacking further theoretical methods for accelerating the convergence of the Robbins-Monro process, methods similar to these seem to be worthy of further study.

6. The block averaging technique holds some promise but the problems involving the choice of an initial c, the method of decreasing it and a "suitable" block size are apparent and need to be studied. Also, other versions of the idea deserve consideration. So far the method seems to

to have the effect of stabilizing the x_n in the vicinity of θ . We feel this is the best method where large sample sizes are not available and the experimenter is lacking prior data.

7. Except, possibly, for a technique such as block averaging, the value of c does not seem to be as critical as it might appear to be at first glance. The experimenter should, of course, know the approximate magnitudes of the variables near level α and can adjust c accordingly. Depending upon these magnitudes, "moderate" values for c should be chosen, assuming a good guess, to keep x_n from straying too far away from θ on the first few iterations.

As noted previously, for the N(0,1), $1 < c < 5$ seems to be the best range for c when the guesses are within 0.5 of the true θ point. We feel justified in assuming this would also be the case for a large class of distributions possessing tails within a fairly wide region about that of the N(0,1). For more general distributions, it seems appropriate to measure c in more general units, say standard deviations σ or reciprocal slope $1/\alpha_1$.

E. Recommendations for Further Study

1. The value of information obtainable by using a method similar to the pseudo confidence interval has been mentioned in conclusion 5. It is believed they are worthy of further study not only computationally, in an effort to obtain knowledge of the approximate sample sizes required to set bounds on the $\{x_n\}$ for different distributions, but analytically, in an effort to establish the degree of confidence one has in placing the bounds when using this method.

In other words, some hypothesis testing scheme might be studied which would test the hypothesis that, when a bound is set on the $\{x_n\}$, it is on the correct side of Θ.

2. Some method for damping the effect of a series of undesirable events, given the guess, x_o, was close to Θ, in order to prevent $\{x_n\}$ from straying too far away from Θ might be studied.

A particular method in this class is the Cesaro sum method of averaging. That is, it is known that if a sequence $\{x_n\} \rightarrow \Theta$ then the sequence of averages $\{x_n^*\} \rightarrow \Theta$ where $x_n^* = 1/n \sum_{i=1}^{n} x_i$. This technique is used to stabilize oscillating sequences and may result in accelerating the convergence for "close" guesses.

3. Kesten (6) develops a method for accelerating convergence of the process based on the number of times the difference of the x_n changes in sign. That is, he reasons that the closer to Θ the x_n gets the more frequently the sign of their difference should change. This method should be studied further from a computational point of view and it may prove helpful given the x_n actually are "close enough" to Θ.

4. For the interested reader the various methods discussed in this paper and those recommended above should, of course, be studied using

widely differing distributions to determine the class of distributions in which these methods could be practically used in accelerating convergence.

BIBLIOGRAPHY

1. Robbins, H., Monro, S., "A Stochastic Approximation Method", Annals of Mathematical Statistics, Vol. 22, 1951, pp. 400-407.

2. Blum, J. R., "Approximation Methods which Converge with Probability One", Annals of Mathematical Statistics, Vol. 25, 1954, pp. 382-386.

3. Chung, K.L., "On a Stochastic Approximation Method", Annals of Mathematical Statistics, Vol. 25, 1954, pp. 463-483.

4. Hodges, J.L., Jr., Lehmann, E.L., "Two Approximations to the Robbins-Monro Process", Proceedings of the Third Berkeley Symposium, Vol. 1, University of California Press, 1956, pp. 96-104.

5. Finney, D.J., "Probit Analysis", Cambridge University Press, 1952.

6. Kesten, H., "Accelerated Stochastic Approximation", Annals of Mathematical Statistics, Vol. 29, 1958, pp. 41-59.

APPENDIX 1

Investigating the range of c for two distributions by a comparison of the x_n with the true θ point of the distribution, a maximum of 500 iterations was allowed where L is the given comparison value, i.e., $|x_n - \theta| \leq L$ for a cessation of the trials. Unlisted values for any particular c means convergence of the x_n to within L of the θ point was not obtained in 500 iterations. More than one value for n indicates the computation was carried out using the same distribution and initial information but with a different set of random numbers from the distribution.

U(0,1)

($\alpha = \theta$)

1. L = 0.001, α = 0.95:

x_o	c	n
0.9	0.5	37
	0.8	24
	1.0	31
0.8	0.6	67,19
	0.7	182
	0.8	27,68
	1.0	16,39
0.7	0.5	16
	0.9	14
0.6	0.4	101
	0.5	238
	1.0	12
0.5	0.6	67
	0.8	38

Note:
c = 0.1(0.1)1.0
only those cases which "converged" under n = 500 are listed

2. L = 0.001, x_o = 0.8

α	c	n
0.96	0.5	49
	0.7	26
	0.8	16
	1.0	90
0.97	0.9	52
0.98	0.4	7
	0.7	352
	0.8	103
	1.0	19
0.99	none	-

3. α = 0.95, x_o = 0.8

L	c	n
0.002	0.5	210
	0.7	16
	0.8	11
	1.0	78
0.003	0.5	28
	0.8	246
	0.9	13
0.004	0.6	17
	0.7	178
	1.0	36
0.005	0.6	142
	0.9	157
	1.0	5

N(0,1)

$\alpha = 0.95(\Theta = 1.645)$

c	x_o	n
1.6	1.0	53
1.7	1.2	16
1.7	2.0	39
1.8	1.5	7
1.8	2.0	31
1.9	2.0	80
1.95	2.0	23
2.55	1.0	415
2.55	2.0	10
3.05	2.0	172
3.55	1.5	147
3.55	2.5	321
4.05	1.2	144
4.05	1.5	357
4.05	2.0	233
4.55	1.2	78
4.55	1.5	45
4.55	2.5	25
5.55	2.5	115
6.05	2.5	43
6.55	1.2	116
6.55	2.5	139

Note: c = 0.2 - 6.55, x_o = 1.0(0.1)2.5

only those cases which "converged" under

n = 500 are listed.

APPENDIX 2

$\alpha = 0.95$, $F^{-1}(\alpha) = 1.645$

200 iterations allowed

L = length of allowed interval = 0.4

A = lower barrier, B = upper barrier

x_o	A	B	c	x_{200}	Lower Reflections	Upper Reflections
1.2	1.1	1.5	5	1.4035	20	1
1.2	1.2	1.6	5	1.5794	18	6
1.3	1.1	1.5	5	1.4076	19	1 (at x_2)
1.3	1.1	1.5	10	1.36	20	5
1.4	1.2	1.6	5	1.5733	16	5
1.4	1.2	1.6	10	1.55	20	6
1.5	1.3	1.7	5	1.6725	9	3
1.5	1.3	1.7	10	1.63	17	7
1.6	1.4	1.8	5	1.5433	19	1
1.6	1.4	1.8	10	1.55	20	1
1.7	1.5	1.9	2	1.7459	2	0
1.7	1.5	1.9	5	1.5865	77	0
1.7	1.5	1.9	10	1.67	78	0
1.8	1.6	2.0	5	1.8427	45	2
1.8	1.6	2.0	10	1.93	78	4
1.9	1.7	2.1	5	1.7081	14	0
1.9	1.7	2.1	10	1.70	49	2

x_o	L	A	B	c	x_{200}	Lower Reflections	Upper Reflections
1.4	0.8	1.0	1.8	2	1.6652	0	1
1.4	1.0	0.9	1.9	2	1.6943	0	0
1.7	0.6	1.4	2.0	2	1.7085	0	0
1.8	0.6	1.5	2.1	2	1.550	12	0
1.8	0.8	1.4	2.2	2	1.7026	0	0

APPENDIX 3

Sample mean and sample variance of m repetitions of the process

$$\bar{x}_i = \frac{1}{m} \sum_{i=1}^{m} x_{ij} \qquad s_i^2 = \frac{1}{m} \sum_{i-1}^{m} (x_{ij} - \bar{x}_i)^2$$

Note that s_i^2 falls off as the logarithm of i

Distribution curve is

Normal (0,1)

$\alpha = .95$

$\Theta = 1.645$

$x_o = 2.5$

$c = 10.0$

$m = 25$

c = 10	i	$\bar{x}_n$	s_i^2	$\frac{1}{i} \frac{\sigma^2 c^2}{2\alpha_1 - 1}$
	25	2.167	1.410	4.44
	50	2.137	1.081	2.22
	75	2.060	0.972	1.49
	100	2.032	0.861	1.11
	125	2.012	0.814	.980
	150	2.011	0.747	.749
	175	2.009	0.699	.632
	200	2.006	0.659	.556
	225	1.996	0.626	.492
	250	1.989	0.602	.486
	275	1.991	0.569	.402
	300	1.978	0.533	.362

APPENDIX 4

Each block is 20 iterations. The current block average provides the "guess" for the next block. The constant c is halved at the beginning of each block.

N(0,1)

$\alpha = 0.950$

$\theta = 1.645$

x_o	n	c	$x_{m,n}$	x_{ave}
1.8	20	5.0	1.213	1.213
	40	2.5	1.932	1.572
	60	1.25	1.588	1.652
	80	0.625	1.750	1.676
	100	0.3125	1.722	1.686
	120	0.1563	1.709	1.689
	140	0.0781	1.709	1.692
	160	0.0391	1.707	1.694
	180	0.0195	1.706	1.695
	200	0.0098	1.705	1.696
1.5	20	10	1.701	1.701
	40	5.0	2.073	1.887
	60	2.5	1.831	1.868
	80	1.25	1.710	1.829
	100	0.625	1.655	1.794
	120	0.3125	1.617	1.765
	140	0.1563	1.627	1.745
	160	0.0781	1.627	1.730
	180	0.0391	1.625	1.719
	200	0.0195	1.621	1.709

$c = 1_n$ 20, ln ln 20, ln ln ln 20, etc.

x_o	I	c	x	x_{ave}
1.8	20	2.9957	1.636	1.636
	40	1.0972	1.638	1.637
	60	0.0928	1.630	1.633

APPENDIX 5

Nesting the Binomial

For computational purposes the binomial sum may be nested by

$$\sum_{i=x}^{n} \binom{n}{i} p^i q^{n-i} = p^n \left(1 + n \frac{q}{p} \left(1 + \frac{n-1}{2} \frac{q}{p} \left(1 + \frac{n-2}{3} \frac{q}{p} \left(\ldots\right.\right.\right.$$

$$\left.\left.\left.\ldots\left(1 + \frac{x+2}{n-(x+1)} \frac{q}{p} \left(1 + \frac{x+1}{n-x} \frac{q}{p}\right)\ldots\right)\right.\right.\right.$$

$$= p^n\left(1 + n \frac{q}{p} \left(1 + \frac{n-1}{2} \frac{q}{p} \left(1 + \frac{n-2}{3} \frac{q}{p} \left(\ldots\left(1 + \frac{j+1}{n-j} \frac{q}{p} \left(1 + \frac{j}{n-(j-1)} \frac{q}{p}\right)\ldots\right)\right.\right.\right.\right.$$

where $j = x + 1, \ldots, n$.

Let $A_j = \frac{j}{n-(j-1)} \frac{q}{p} \Longrightarrow$

$$\sum_{i=x}^{n} \binom{n}{i} p^i q^{n-1} = p^n(1 + \ldots(1 + A_{j+2} (1 + A_{j+1} (1 + A_j)\ldots)$$

Let $B_1 = 1$ and $B_{k+1} = 1 + A_j B_k$, $k = 1, \ldots, n-x+1$

where n-x+1 is the number of terms in the series expansion of the binomial sum.

$$\Longrightarrow \quad B_2 = 1 + A_{x+1}\, B_1 = 1 + \frac{x+1}{n-x}\,\frac{q}{p}$$

$$B_3 = 1 + A_{x+2}\, B_2 = 1 + \left(\frac{x+2}{n-(x+1)}\,\frac{q}{p}\right)\left(1 + \frac{x+1}{n-x}\,\frac{q}{p}\right)$$

$$\vdots$$

$$B_{n-x} = 1 + A_{n-1}\, B_{n-x-1} = 1 + \frac{n-1}{2}\, B_{n-x-1}$$

$$B_{n-x+1} = 1 + A_n\, B_{n-x} = 1 + \frac{n}{1}\,\frac{q}{p}\, B_{n-x}$$

$$\therefore \quad p^n\, B_{n-x+1} = \sum_{i=x}^{n} \binom{n}{i}\, p^i q^{n-i}$$

APPENDIX 6A

Sample data using pseudo confidence intervals

F(x) is N(0,1) , $1 - \alpha_2 = .90$

n = Sample size

x_n = Terminal value after N iterations

$\overline{P}$ = Upper "confidence" bound

$\underline{P}$ = Lower "confidence" bound

Reflections high: $\overline{P} < \alpha$

Reflections low: $\underline{P} > \alpha$

$\alpha = .95$ $x_o = 1.4$

c	n	x_n	$\overline{P} - \underline{P}$	Reflections High	Low
2.0	76	1.344	.030	0	0
2.0	143	1.446	.050	0	7
2.0	175	1.450	.054	0	5

$\alpha = .95$ $x_o = 1.5$

c	n	x_n	$\overline{P} - \underline{P}$	Reflections High	Low
2.0	84	1.436	.129	0	0
1.0	84	1.460	.129	0	0
ln(n)	88	1.543	.099	0	0
ln(n+1)	94	1.522	.098	0	0
2.0	176	1.500	.082	0	4
ln(n+1)	220	1.691	.038	0	0

$\alpha = .95$ $x_o = 1.6$

c	n	x_n	$\overline{P} - \underline{P}$	Reflections High	Low
1.0	84	1.568	.029	0	0
2.0	143	1.597	.050	0	6
2.0	176	1.600	.062	0	4

α = .95 x_o = 1.7

c	n	x_n	$\bar{P} - \underline{P}$	Reflections High	Low
2.0	86	1.619	.096	0	0
1.0	104	1.644	.110	0	0
2.0	143	1.446	.050	0	0

α = .95 x_o = 1.8

c	n	x_n	$\bar{P} - \underline{P}$	Reflections High	Low
2.0	104	1.688	.114	0	0
2.0	143	1.707	.050	0	3

α = .95 x_o = 1.9

c	n	x_n	$\bar{P} - \underline{P}$	Reflections High	Low
2.0	116	1.677	.094	0	0
2.0	143	1.689	.050	0	2
2.0	176	1.692	.062	0	0

α = .95 x_o = 2.0

c	n	x_n	$\bar{P} - \underline{P}$	Reflections High	Low
2.0	28	1.607	.098	0	0
2.0	102	1.808	.049	0	1
4.0	121	1.733	.098	0	0
2.0	122	1.790	.051	0	1
10.0	132	2.018	.089	0	0
1.0	167	1.831	.050	0	0
2.0	200	1.782	.052	0	0
2.0	332	1.686	.049	0	1
ln(n+1)	425	1.598	.037	0	0
1.0	537	1.740	.048	0	0
2.0	579	1.796	.044	0	0

α = .95 x_o = 2.2

c	n	x_n	$\bar{P} - \underline{P}$	Reflections High	Low
2.0	156	1.864	.015	0	0
10.0	248	2.00	.042	0	1

APPENDIX 6B

$\alpha = .95, \quad \theta = 1.645$

$x_o = 2.0, \quad \alpha_2 = .1, \quad L = .02$

	Successes	Failures	x_n	$\bar{P}$	$\underline{P}$
C=1:	57		1.76855	1.0	.94966
	$\underline{P}$ tried to to high		changed success to failure		
		1*	1.78493	.99913	.92213
	24		1.76774	.99937	.94343
		1	1.77919	.99570	.92611
	28		1.76473	.99679	.94435
		1	1.77321	.99274	.93251
	8		1.76977	.99322	.93680
		1	1.77763	.98877	.92630
	58		1.75811	.99250	.94998
	$\underline{P}$ tried to go high				
		1*	1.76339	.98930	.94306
	12		1.76017	.98996	.94622
		2	1.76999	.98346	.93387
	little change in numbers				
	total # times $\underline{P}$ tried to go high = 4				
			x_{1105} = 1.75674	.98470	.94753
				$\bar{P} - \underline{P}$ =	.03717

Note: * = deliberately changed actual experimental result

	Successes	Failures	x_n	$\bar{P}$	$\underline{P}$
C=2:	48		1.55412	1.0	.93950
		1	1.59290	.99895	.90685
	33		1.54181	.99937	.94343
		1	1.56471	.99570	.92611
	28		1.53579	.99679	.94435
		1	1.55275	.99274	.93251
	8		1.54588	.99322	.93680
		1	1.56159	.98877	.92630
	little change no reflections recorded				
15 min. limit:	x_{1003} =		1.57895	.98650	.93236
				$\bar{P} - \underline{P}$ =	.05414

	Successes	Failures	x_n	$\bar{P}$	$\underline{P}$
C=3:	4		1.68750	1.0	.47287
		1	2.25750	.98979	.34272
	7		2.13452	.99573	.66121
		1	2.35375	.97196	.59001
	68		2.08410	.99559	.92429
		1	2.11885	.99011	.90866
	69		2.02768	.99464	.94983
	$\underline{P}$ tried to go high so changed success to failure				
		1*	2.04643	.99107	.94133
	27		2.02198	.99250	.94998
	$\underline{P}$ tried to go high				
		1*	2.03782	.98930	.94306
	little change in numbers, 15 more bounces recorded where $\underline{P}$ tried to go high and had to change success to failure				
	total # reflections recorded = 17				
		x_{1135} =	1.99758	.98542	.94897
				$\bar{P} - \underline{P}$ =	.03645
C=5:	4		1.47917	1.0	.47287
		1	2.42917	.98979	.34272
	7		2.22420	.99573	.66121
		1	2.58958	.97196	.59001
	68		2.14016	.99559	.92429
		1	2.19809	.99011	.90866
	69		2.04614	.99464	.94983
	$\underline{P}$ tried to go high, changed success to failure				
		1*	2-07739	.99107	.94133
	little change in numbers				
	total # times $\underline{P}$ tried to high = 15				
15 min:		x_{1077}	1.99561	.98762	.94954
				$\bar{P} - \underline{P}$ =	.03808

	Successes	Failures	x_n	$\bar{P}$	$\underline{P}$
C=ln(n+1):	17		1.73875	1.0	.83843
		1	1.89415	.99715	.76243
	31		1.72470	.99895	.90677
		1	1.79941	.99285	.87980
	21		1.72791	.99497	.91394
		1	1.78452	.98875	.89629
	4		1.77291	.98933	.90167
		1	1.82666	.98207	.88542
	7		1.80761	.98417	.89486
		1	1.85739	.97768	.88128
	9		1.83483	.98053	.89242
		1	1.88047	.97226	.87934
	18		1.84028	.97741	.89860
		1	1.87982	.97404	.88882
	little change, $\underline{P}$ tried to go high 1 time				
		x_{1087} =	1.75673	.98739	.94769
				$\bar{P} - \underline{P}$ =	.03970
C=6:	4		1.37500	1.0	.47287
		1	2.51500	.98980	.34390
	7		2.26904	.99573	.66018
		1	2.70750	.97207	.59087
	68		2.16819	.99559	.92406
		1	2.23770	.99014	.90889
	70		2.05339	.99463	.95000
	$\underline{P}$ tried to go high		changed success to failure		
		1*	2.09065	.99133	.94200
	26		2.04371	.99248	.94986
	$\underline{P}$ tried to go high				
		1*	2.07537	.98934	.94319
	26		2.03500	.99053	.94988
	$\underline{P}$ tried to go high				
		1*	2.06254	.98805	.94441
	little change in numbers				
	total # times $\underline{P}$ tried to go high = 13				
		x_{1000} =	1.99340	.98436	.94939
				$\bar{P} - \underline{P}$ =	.03497

	Successes	Failures	x_n	$\bar{P}$	$\underline{P}$
C=1	57		1.76855	1.0	.94966
	$\underline{P}$ tried to go high				
		1*	1.78493	.99913	.92240
	24		1.76774	.99937	.94321
		1	1.77919	.99571	.92635
	28		1.76473	.99678	.94417
		1	1.77277	.99277	.93270
	8		1.76977	.99321	.93664
		1	1.77763	.98880	.92648
	58		1.75811	.99248	.94986
	$\underline{P}$ tried to go high				
		1*	1.76339	.98934	.94319
	total #times $\underline{P}$ tried to go high = 3				
		x_{1000} =	1.75490	.98436	.94939
				$\bar{P} - \underline{P}$ =	.03497

APPENDIX 7

Proof of the convergence of $\frac{(\ln n)^2}{n}$

Let $a_n = K \frac{\log_{10} n}{n}$

$$\sum_{n=1}^{\infty} a_n^2 = \sum_{m=o}^{\infty} \sum_{n=10^m+1}^{10^{m+1}} a_n^2 = K^2 \sum_{m=0}^{\infty} \sum_{n=10^m+1}^{10^{m+1}} \frac{\log^2 n}{n^2}$$

For $10^m+1 \leq n \leq 10^{m+1}$ we have

$$\log^2 n \leq (\log 10^{m+1})^2 = (m+1)^2$$

also $$\sum_{10^m+1}^{10^{m+1}} \frac{1}{n^2} \leq \int_{10^m}^{10^{m+1}} \frac{dx}{x^2} = (-\frac{1}{x})_{10^m}^{10^{m+1}} = \frac{9}{10^m}$$

$$\therefore \sum_{n=1}^{\infty} a_n^2 \leq K^2 \sum_{m=0}^{\infty} \frac{9(m+1)^2}{10^m} < \infty$$

since it is dominated by a geometric series.